

# Using a Model of Ocean Currents to Control the Position of Vertically Profiling Marine Floats

**Martina Troesch, Steve Chien**

Jet Propulsion Laboratory  
California Institute of Technology  
4800 Oak Grove Drive  
Pasadena, CA 91109  
{firstname.lastname}@jpl.nasa.gov

**Yi Chao, John Farrara**

Remote Sensing Solutions  
248 East Foothill Blvd  
Monrovia, CA 91016  
ychao@remotesensingsolutions.com  
jfarrara@remotesensingsolutions.com

## Abstract

We describe a methodology for control of vertically profiling floats that uses an imperfect predictive model of ocean currents. In this approach, the floats have control only over their depth. We combine this control authority with an imperfect model of ocean currents to force the floats to maintain position. First, we study the impact of model accuracy on this ability to station keep (e.g. maintain X-Y position) using simulated planning and nature models. In this study, we examine the impact of batch versus continuous planning. In batch planning the float depth plan is derived for an extended period of time and then executed open loop. In continuous planning the depth plan is updated with the actual position and the remainder of the plan re-planned based on the new information. In these simulation results, we show that (a) active control can significantly improve station keeping with even an imperfect predictive model and (b) continuous planning can mitigate the impact of model inaccuracy. Second, we study the effect of using heuristic path completion estimators in search. In general, using a more conservative estimator increases search quality but commensurately increases the amount of search and therefore computation time. Third, we discuss results from an April 2015 deployment in the Pacific Ocean and compare model accuracy and float control performance.

## Introduction

The state of the ocean affects the environment and climate, thus affecting food production, defense, and leisure. As such, ocean dynamics is an important area of study that currently uses a variety of different techniques to measure ocean conditions. One technique involves the use of robotic marine vehicles such as floats, gliders, and autonomous underwater vehicles (AUV) to measure conditions such as currents, salinity, and temperature in a dynamic way. Another technique uses moored buoys, which allow scientists to collect data at a fixed location over time. However, a couple of drawbacks to physically mooring a buoy are that it involves significant financial investment and the location cannot be changed after installation.

As an alternative, a *virtual mooring* is proposed in which a dynamically controlled vehicle uses a control policy in order to maintain its position. Specifically, one proposal is to

Asset	Control	Speed	Longevity	Cost
Floater	None	None	Weeks	\$100's
Vertical Profiling Float	Vertical	~0.1 m/s	Years	\$10K's
Seaglider	Horizontal	~0.5 m/s	Months	\$50's - \$100'sK
AUV	Horizontal & Vertical	~2.5 m/s	1 Day - Weeks	\$100K - \$M

Table 1: Characteristics and costs for different families of marine vehicles.

deploy a vertical profiling float to the location of desired data collection and to use predictive ocean models to plan a control sequence for changing depths that best keeps the float near the same latitudinal and longitudinal location using the ocean currents. A vertical profiling float can change depths, but does not have any lateral propulsive power, meaning that the float is carried solely by the ocean current in the latitudinal and longitudinal directions. By purposefully changing depths it is possible to harness the motion of the ocean to direct the float. This method has multiple benefits over using a physical mooring. First, the float could be retrieved and redeployed when desired. Second, there is more flexibility since the float could be programmed to track a moving target or to drift to facilitate deployment or retrieval. Third, the deployment would be less expensive than building a physical anchor location.

Although using an AUV would provide better control to remain at a fixed location, more capable vehicles are more expensive. Table 1 shows approximate costs for families of marine vehicles (Woods Hole Oceanographic Institution ; Sanford et al. 2005; Eriksen et al. 2001; YSI Systems ; OceanServer Technology, Inc. ; Bluefin Robotics Corporation ; Kongsberg Mairtime AS ).

Scientists studying the characteristics of the ocean would ideally like to be able to collect data at all depths and all times at a particular location. Obviously, a single float cannot be at all depths at all times and therefore must profile to collect data across the depths. Using a predictive ocean model it is possible to generate a control sequence for the float to change depths in a way that keeps it as close to the desired location of data collection as possible.

To analyze the benefit of planning a path for a float to act as a virtual mooring, compared to allowing the float to continuously profile, an Electromagnetic Autonomous Profiling Explorer (EM-APEX) (Sanford et al. 2005) vertical profiling float is modeled. During a deployment, each time that the EM-APEX float surfaces, it transmits its data and can be commanded to profile to a different depth. This allows for two possible control strategies. First, in *batch planning*, the float plans once for the deployment based on the best model of the ocean currents. In *continuous planning*, at each float surfacing, we re-plan the control sequence of the float during each surfacing. This enables the planning to incorporate: (1) the most up to date information about the location of the float and (2) the most up to date ocean current model.<sup>1</sup> We believe that using this opportunity to re-plan the control sequence using the best information can improve the path when there is information gain in the model.

Since directing the float relies solely on the ocean currents, success is based heavily on the planning process, and thus on the ocean model used for planning. Modeling ocean currents is a tremendously challenging problem - as a chaotic system it is not feasible to model the ocean perfectly and producing even modestly accurate predictions is quite challenging. As we investigate the use of predictive ocean current models to plan underwater vehicle paths, predictive accuracy can dramatically affect the utility of our approach. Therefore studying the impact of model accuracy on path planning is of great import. Additionally, methods of measuring predictive model accuracy and correlating these to planning utility is of great interest.

To study this relationship between model accuracy and planning utility, we use the Regional Ocean Modeling System (ROMS) (Chao et al. 2009; Li et al. 2006; Farrara et al. 2015). Specifically we artificially create models with varying degrees of fidelity. Because it is very expensive to perform a physical deployment in the ocean, we mimic a deployment. In a deployment we plan in a model and we execute in the physical world. In ROMS, we create one or more planning models of varying fidelity to a nature model. We then construct plans in a planning model and execute in the nature model. The planning models were five other ROMS models with decreasing fidelity. In order to evaluate the models, we use the models for path planning of an Electromagnetic Autonomous Profiling Explorer (EM-APEX) (Sanford et al. 2005) vertical profiling float attempting to act as a virtual mooring.

The paths were planned in all of the models and compared to the execution in the nature model. The paths were also re-planned using continuous planning during execution using each of the models to compare to the results without re-planning. We show that using current model to plan a path for a vertical profiling float to act as a virtual mooring can improve its station keeping compared to a naive control strategy and that re-planning the path during execution

is an effective technique. Furthermore, this paper aims to show how the information in a model affects the benefit of planning a path as well as the efficacy of re-planning during execution. These ideas have also previously been explored in (Troesch et al. 2016).

The remainder of this paper is organized as follows. First, we describe the ROMS modeling framework that we use as an imperfect predictive model of ocean currents. Second, we describe the batch and continuous float planning algorithms. Third, we describe our results in simulation - highlighting both the effect of model accuracy and heuristic path completion estimation on algorithm performance. Fourth, we describe results from an April 2015 field deployment off the coast of California. Fifth, we describe related and future work, and conclusions.

## Ocean Models

A number of ocean models have been developed including the Harvard Ocean Prediction System (HOPS) (Robinson 1999), the Princeton Ocean Model (POM) (Mellor 2004), the Hybrid Coordinate Ocean Model (HYCOM) (Chassignet et al. 2007), and ROMS (Chao et al. 2009). As described in the Introduction, for our experiments we use the ROMS model. However, our techniques naturally extend to any cell-based, predictive model with information about ocean currents over multiple depths and an extended period of time, and thus any of these models could be used for the path planning. Indeed, when multiple models are available it is also possible to use them in as an ensemble to further enhance results (Wang et al. 2013).

ROMS is a discrete, cell-based, predictive model of the ocean. We used the California coast configuration near the Monterey Bay, which is a grid of 3 km by 3 km in the latitudinal and longitudinal directions, 1 hour in the time dimension up to 72 hours long, and fourteen depths from 0 m to 1000 m in non-uniform intervals. The currents in the grid vary over space and time. At deeper depths, the currents tend to be slow and uniform, conversely, the surface currents are faster and more variable.

As previously stated, it is not feasible to perform an ocean deployment of the planned paths for this experiment, so an approximation using a ROMS model for the ocean is used instead. This model is the best possible representation from ROMS and is referred to as the nature model. Five different planning models were used for this experiment. The difference between the models is the number of days advanced prediction that is used in generating the model. Fewer days of advanced prediction means a higher fidelity model and thus means that the model is closer to the nature model. We used 2, 4, 6, 8, and 10 days of advanced prediction for the planning models. A summary of the inputs for the ROMS models is shown in Table 2.

To show how the model information decreases with more advanced prediction of the model, the correlation coefficient of the currents between the nature model and each of the planning models was calculated. All of the velocities in a 41 grid by 41 grid subsection of the ROMS model across all depths and times were used for the calculation. In

<sup>1</sup>While in our deployment scenarios the model does not change significantly on the timescale of our plan execution so that majority of the gain is from (1), other operational scenarios might exist where (2) may provide significant value

	Planning Models	Nature Model
Archiving, Validation and Interpretation of Satellite Oceanographic (AVISO) sea surface height data	x	x
Advanced Very High Resolution Radiometer (AVHRR) sea surface temperatures	x	x
Moderate Resolution Infrared Spectroradiometer (MODIS) sea surface temperatures	x	x
GOES satellite sea surface temperatures	x	x
High Frequency (HF) radar surface current data		x
Monterey Bay Aquarium Research Institute (MBARI) M1 mooring vertical profiles of temperature and salinity	x	x
Ship sea surface temperatures	x	x
Number of days advanced prediction	2, 4, 6, 8, 10	1

Table 2: ROMS inputs for the planning and nature models.




Figure 1: The correlation coefficients of the zonal and meridional currents between the different planning models and the nature model over a 41 by 41 3km x 3km cell subsection of the model averaged over all depths and times.

other words,  $Z_{nature}$  is a vector of the zonal (west-east) currents in the selected subregion of the nature ROMS model at all depths and times. The vector  $M_{nature}$  contains the corresponding meridional (north-south) currents. Similarly,  $Z_{plan_x}$  and  $M_{plan_x}$  contain the zonal and meridional cur-




Figure 2: The root mean square error of the zonal and meridional currents between the different planning models and the nature model over a grid that encompasses the entire search space over all locations at all depths and times.

Days Advanced Prediction	RMS Error
2	0.00855
4	0.05665
6	0.08555
8	0.11445
10	0.13735

Table 3: Root mean square error used for each model with the specified number of days advanced prediction.

rents of the same data in the planning model using  $x$  days of advanced prediction, respectively. Using these vectors, the zonal correlation coefficients,  $\rho_{Z_x}$ , and the meridional correlation coefficients,  $\rho_{M_x}$ , could be calculated with the following equations

$$\rho_{Z_x} = \frac{\text{Cov}(Z_{plan_x}, Z_{nature})}{\sigma_{Z_{plan_x}} \sigma_{Z_{nature}}}$$

$$\rho_{M_x} = \frac{\text{Cov}(M_{plan_x}, M_{nature})}{\sigma_{M_{plan_x}} \sigma_{M_{nature}}}$$

where Cov is the covariance function and  $\sigma_{Z_{plan_x}}$ ,  $\sigma_{Z_{nature}}$ ,  $\sigma_{M_{plan_x}}$ , and  $\sigma_{M_{nature}}$  are the standard deviations of the referenced vector.

Figure 1 summarizes the values of the correlation coefficients and shows how the information in the model becomes less correlated to the nature model as the number of advanced prediction days increases. In fact, the correlation coefficient for the zonal currents using 8 and 10 days of advanced prediction has a negative value, indicating that the




Figure 3: The currents at a surface point over 72 hours in all of the planning models and the nature model.

currents are inversely correlated to the currents in the nature model, which causes the value of planning a path to decrease.

Just as the correlation coefficient decreases with the increase in number of advanced prediction days, the root mean square (RMS) error increases. Figure 2 shows the RMS error between each planning model and the nature model using the currents over the search space that encompasses all of the goal locations using all depths and times. The average RMS error between the zonal and meridional currents for each of the planning models will be used in the rest of this paper as a measure of the fidelity of the model, which in turn could be used as a comparison and reference point in physical deployments to analyze the fidelity of the model being used. The values of RMS error used for each model are shown in Table 3.

Another way to visualize how the currents vary across the different models is by looking at the currents at one location over time across all models. Figure 3 shows the zonal and meridional currents for the same location in all models over time. The currents of the planning model with 2 days of prediction are very close to the nature model, but the other planning models diverge.

Furthermore, Fig. 4 demonstrates how the exact same control sequence executed over 24 hours in the different models results in very different path positions.

From Fig. 1, Fig. 2, Fig. 3, and Fig. 4 it is clear that there is different value in the information provided by the different planning models. Specifically that with further advanced prediction there is a decrease in the predictive accuracy of the model as indicated by the decrease in correlation coefficient and increase in RMS error. Furthermore, that these

varying information models will predict significantly varying paths and consequently planing in the different models can produce significantly different paths.



Figure 4: The same executed control sequence in all of the planning models and the nature model starting at the same location over 24 hours. The number 0 labels the path for the nature model and the other numbers indicate the number of planning days in the model used for that path.

## Batch Planning

The path planner for this experiment generates a control sequence for a vertical profiling float to act as a virtual mooring. As described in the introduction, a vertical profiling float can change depths both to gather data and to make use of the currents to stay near a desired location. The float can be programmed to move between the surface and the profiling depth, to remain at the surface, or to remain at the profiling depth. Stopping and staying at a depth is restricted to the surface and profiling depth in order to best mimic the behavior of the EM-APEX floats. Depending on the needs of the scientist, it may be more important to stay at a fixed location or to gather more data by profiling. In this way a trade-off can be made between staying at the desired location and performing more profiles.

Even though the problem space is continuous, to make the search tractable, the control sequence of the float is determined in a discrete manner. At each decision point, the float can remain at the depth that it is at, which must be

the surface or the profiling depth, or it can move between the surface and the profiling depth. If the float remains at a depth, the duration is equivalent to the time it takes to move between the depths. To mimic the behavior of a deployed EM-APEX float, each time that the float returns to the surface from the profiling depth, it must remain at the surface to upload the collected data. The duration of the upload is 35 minutes. If the float decides to remain at the surface, it must once again upload the collected data for 35 minutes after the duration of remaining at the surface.

The value of the currents used for determining the motion of the float is based on the position of the float as well as the time. Every approximately 42 seconds (the amount of time that it takes the float to vertically move half of the smallest depth step), the position is updated and the current information is interpolated among the eight closest grid points. This current is used to determine the motion of the float until the next interpolation step or the allotted time step for the node has been reached.

The algorithm that is used to perform the search is an A\* algorithm. The objective function that is used to score the paths was designed to make a trade-off between remaining close to a desired location and performing more profiles, depending on the desires of the user. The equation is

$$g(n) = \sum_n \sum_d (w_T \Delta T_d + w_D \Delta D_d)$$

where  $n$  are the nodes in the path,  $d$  are the possible depth choices (in this case, the surface or the profiling depth),  $w_T$  and  $w_D$  are weighting terms,  $\Delta T_d$  is the time in seconds since the float was at depth  $d$ , and  $\Delta D$  is the distance in kilometers that the float was from the goal location when it was last at depth  $d$ . In other words, each time a node is added to the path, the most recent node at the surface and the profiling depth is used to calculate  $(w_T \Delta T_d + w_D \Delta D_d)$  and this sum is added to the score. Since  $\Delta T_d$  measures the time since the float was at the other depth, it is a good metric for determining the time since the last profile was performed. Therefore, a higher  $w_T$  to  $w_D$  ratio favors less time between profiles and thus favors more profiles. Similarly, a lower  $w_T$  to  $w_D$  ratio favors control sequences that keep the float as close as possible to the desired location.

The heuristic function used in the A\* algorithm simply assumes that the score will increase at the same rate for the remainder of the path. Therefore, the equation is

$$h(n) = \left( \frac{g(n)}{T} \right) (\text{mission duration} - T)$$

where  $T$  is the time since the beginning of the mission when the float is at node  $n$ .

The equation used by the A\* algorithm is thus

$$f(n) = g(n) + w * h(n)$$

where  $w$  is a weighting given to the heuristic function.

A path is considered complete once its duration reaches that of the desired mission length.

The execution of the algorithm is summarized in the following pseudocode for Algorithm 1.

---

#### Algorithm 1 Batch Planning A\* Algorithm

---

```

1: function BATCHPLANNER(startPath)
2: (*Note: Model = Planning Model)
3:    $Q \leftarrow \text{startPath}$ 
4:   while  $Q$  not empty do
5:      $\text{curPath} \leftarrow$  lowest  $f$  score path in  $Q$ 
6:     if  $\text{curPath}$  needs to upload then
7:        $\text{newPath0} \leftarrow \text{curPath} + \text{node at surface}$ 
8:     else
9:        $\text{newPath1} \leftarrow \text{curPath} + \text{node at surface}$ 
10:       $\text{newPath2} \leftarrow \text{curPath} + \text{node at profiling}$ 
11:      depth
12:    end if
13:    if any  $\text{newPath}$  duration > mission duration
14:    then
15:      return  $\text{curPath}$ 
16:    end if
17:     $Q.\text{push all newPaths}$ 
18:  end while
19: end function

```

---

### Continuous Planning

During every data upload when the float resurfaces, there is an opportunity to re-plan the path of the float based on the best information of its location. Each time the float performs an upload, the location at the true position of the path so far in the execution can be used in Algorithm 1 to plan the rest of the path using the planning model. The next part of the control sequence is executed according to the re-planned path. The process of re-planning is repeated each time that the float re-surfaces until the duration of the mission has been completed. The control sequence is always executed in the nature model and planned in the planning model. A summary of the re-planning algorithm is shown in Algorithm 2.

---

#### Algorithm 2 Continuous Planning Algorithm

---

```

1: function CONTINUOUSPLANNER(startNode)
2: (*Note: Model = Nature Model)
3:    $\text{path} \leftarrow \text{startNode}$ 
4:    $\text{controlSeq} \leftarrow \text{BATCHPLANNER}(\text{path})$ 
5:   while  $\text{controlSeq}$  duration >  $\text{path}$  duration do
6:     while  $\text{path}$  does not need to upload do
7:        $\text{nextNode} \leftarrow \text{controlSequence.next}$ 
8:        $\text{depth} \leftarrow \text{depth of nextNode}$ 
9:        $\text{path} \leftarrow \text{path} + \text{node at depth}$ 
10:    end while
11:     $\text{controlSeq} \leftarrow \text{BATCHPLANNER}(\text{path})$ 
12:  end while
13:  return  $\text{path}$ 
14: end function

```

---

### Experimental Procedure

The experiment was performed over 100 goal locations in a 10 by 10 grid, which are shown in Fig. 5. At each location, the five different planning models were used to plan a control sequence using Algorithm 1, which were then executed



in the nature model. Each of these control sequences were then also re-planned with continuous planning using Algorithm 2. This process was repeated over all of the locations using 3 different trade-offs in the objective function, specifically a  $w_T$  to  $w_D$  ratio of 1.6, 4.8, and 8.0 were used, as well as 6 different weights for the heuristic function from 0 to 1 in steps of 0.2.

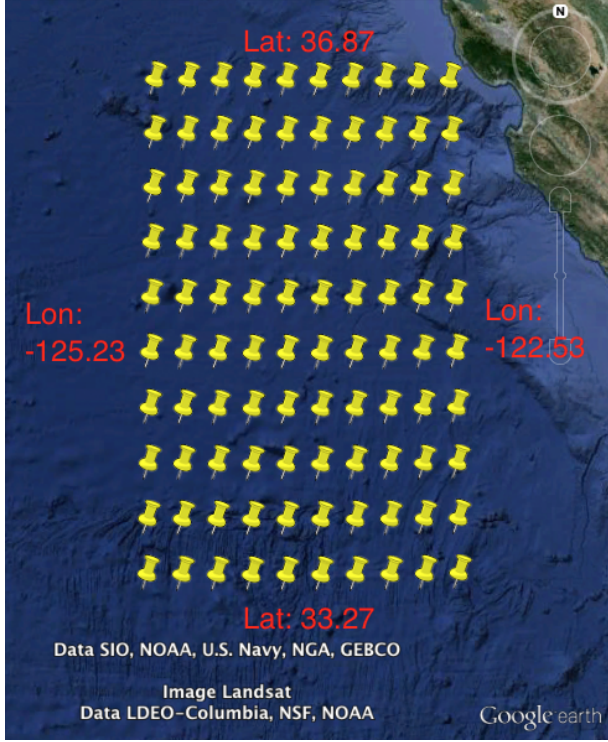


Figure 5: The 100 starting and goal locations used for the experiment. The latitudes and longitudes labeled indicate the boundaries of the grid.

The start location of the float was always set to be the same as the goal location. The profiling depth was 500 m and the vertical speed of the float was 0.12 m/s, resulting in a step time of approximately 69 minutes. As previously stated, the time required for the data upload was 35 minutes. The total duration of the mission was 24 hours. Given these inputs, the maximum number of profiles that can be achieved by continuously profiling is 8.

In order to evaluate the results of the planned paths, baselines were developed that evenly space the profiles, from 0 to the maximum of 8 within the mission duration. In order to attempt to prevent bias in the baselines, three different ways of evenly spacing the profiles were used. The first, referred to as the surface baseline, evenly spaces the profiles by remaining at the surface between profiles that consist of going to the profiling depth and resurfacing immediately. The second baseline, referred to as the even baseline, evenly spaces the profiles by remaining at the surface and at the bottom of the profiles at even intervals. The final baseline, referred to as the deep baseline, remains at the profiling depth dur-




Figure 6: The control sequences for 3 profiles using the surface, even, and deep baselines.

ing each profile and only surfaces in between profiles for the amount of time required to upload the data. As an example, the depths of the different baseline paths for 3 profiles is shown in Fig. 6.

## Evaluation of the Heuristic

In order to evaluate the heuristic function and the effect of the chosen weight, an analysis was done on all of the paths in the planning models. Since the scores are heavily dependent on the  $w_T$  to  $w_D$  ratio in the objective function, the analysis was split among the three executed ratios. The average score and number of node expansions to find a complete path across all locations and planning models was found for each  $w_T$  to  $w_D$  ratio, which can be seen in Fig. 7 and Fig. 8.

As more weight is given to the heuristic function, the strength of the path decreases, which can be seen by the increasing scores in Fig. 7. However, at the same time, the number of nodes expanded decreases, which can be seen in Fig. 8, meaning that the time to find a complete path is improved.

## Empirical Evaluation in Simulation

Using a heuristic weight of 0 in order to ensure using the best planned path, for each  $w_T$  to  $w_D$  ratio, each ROMS model, and each location, the executed path in the nature model was found for both batch planning and continuous planning. As an example to show how the control sequence and the score changes from batch planning to continuous planning, Fig. 9 shows the depth and score over time using 6 days of advanced planning and a  $w_T$  to  $w_D$  ratio of 4.8 at location (35.67° lat, -123.73° lon).

The average scores across all of the locations for each  $w_T$  to  $w_D$  ratio and each ROMS model were then calculated for




Figure 7: Average score across all planning models and locations for different heuristic weights.




Figure 8: Average number of node expansions required across all planning models and locations for different heuristic weights.

both batch planning and continuous planning, separately. In order to give a fair comparison to the baselines, the average number of profiles for each  $w_T$  to  $w_D$  ratio and ROMS model combination was found from both batch planning and continuous planning. This number of profiles was then used

for the baselines with those same inputs. Since the average number of profiles for each  $w_T$  to  $w_D$  ratio was not the same for each ROMS model, the baseline scores are not constant across a single ratio. Comparing the baselines to the different planning methods revealed that the surface baseline performed too poorly (with scores over 100,000) to be represented on the same scale as the other results, therefore the surface baselines are not presented on the graphs showing the results. The comparison of the scores using the different planning models can be seen in Fig. 10 for the  $w_T$  to  $w_D$  ratio of 1.6, in Fig. 11 for the  $w_T$  to  $w_D$  ratio of 4.8, and in Fig. 12 for the  $w_T$  to  $w_D$  ratio of 8.0.




Figure 9: Batch and continuous planning using 6 days of advanced planning and a  $w_T$  to  $w_D$  ratio of 4.8 at location (35.67° lat, -123.73° lon).

From these figures, it is clear that on average planning a control sequence using any of the models performs better than simply using evenly spaced out profiles. Furthermore, they all show that the benefits of planning are increased when the planning model has better knowledge of the currents in the nature model. As the RMS error in the planning models increases, the benefits of planning a control sequence over the baseline decreases.

When considering continuous planning, the benefit is also related to the amount of knowledge in the planning model. Obviously, when planning using the nature model, there is no increase in benefit from continuous planning, since the position was already correctly known during the planning process and thus the continuous plan and the batch plan are identical and there is no added benefit to continuous planning. However, as the planning model and the nature model diverge, the updated information at each surfacing provides a benefit to the planner.

Looking at Fig. 13, the planning model with just a value




Figure 10: The average batch and continuous planning scores when the  $w_T$  to  $w_D$  ratio was 1.6 using each planning model compared to the average baseline scores of the same average number of profiles. The RMS errors for the zonal and meridional currents for two deployed floats discussed in the next section are also shown.




Figure 11: The average batch and continuous planning scores when the  $w_T$  to  $w_D$  ratio was 4.8 using each planning model compared to the average baseline scores of the same average number of profiles. The RMS errors for the zonal and meridional currents for two deployed floats discussed in the next section are also shown.




Figure 12: The average batch and continuous planning scores when the  $w_T$  to  $w_D$  ratio was 8.0 using each planning model compared to the average baseline scores of the same average number of profiles. The RMS errors for the zonal and meridional currents for two deployed floats discussed in the next section are also shown.




Figure 13: The average batch and continuous planning scores when the  $w_T$  to  $w_D$  ratio was 8.0 using each planning model compared to the average baseline scores of the same average number of profiles. The RMS errors for the zonal and meridional currents for two deployed floats discussed in the next section are also shown.



of 0.0085 RMS error is very similar to the nature model, therefore although there is a benefit to continuous planning, the added benefit is not significant since the executed path would not have strayed far from the planned path. Looking at the next three higher values of RMS error, continuous planning has a consistent advantage over batch planning. With the models with the highest values of RMS error, the benefit of continuous planning decreases. This could be due to the fact that starting at such a large value of RMS error, the model has such poor information gain that even though the location is updated at every surfacing, the model is not good enough to make a significantly better plan.

## Empirical Results in Deployment

A prior version of this software was deployed to control EM-APEX floats during an April 2015 deployment in support of an AirSWOT (Jet Propulsion Laboratory a; b) field experiment in the coast off of Monterey Bay, California. In this field experiment, the goal was to keep EM-APEX floats near features of interest identified manually by scientists. The overall AirSWOT deployment goals are representative of the intended scientific use case for these planning tools.

The overall AirSWOT deployment was to test out an Airborne science instrument by providing corroborative data over interesting science features using in-situ instrumentation (floats, ships) and remote sensing data (from overflying spacecraft). The AirSWOT instruments were scheduled to fly in a coverage pattern over specific areas chosen to overlap satellite overflights.

Three EM-APEX floats were to be deployed to be near satellite overflights and airborne overflights. The float planning tool was used to evaluate potential deployment locations by predicting the projected drift path of the floats.

Figure 14 shows the variability of the expected float drift based on the deployment location. The blue paddle indicates the start location and the green path shows where the float was projected to drift over time.

Sites for each of the three float deployments were screened for projected stability and hand selected by the experiment team.

Additionally, two of the three EM-APEX floats were allowed to be controlled dynamically from shore by an earlier version of the float planning software. This prior float planning software received the satellite phone updated location each time the target float surfaced. Because of connectivity issues, the float planning software could not receive this data rapidly enough to generate a new plan for transmission to the float during this surface cycle as the float was only on the surface approximately 30 minutes each cycle. Instead the planner could only assert a plan with a 1 surface cycle lag. Therefore the plan communicated to the float to be executed after surface cycle  $n$  was only based on the actual position from cycle  $n - 1$  plus the projected drift from cycle  $n - 1$  to cycle  $n$ .

The EM-APEX float tracks planned and executed are shown for floats 6665 and 6667 in Fig. 15. The yellow point indicates the start location of the float. The actual location of




Figure 14: Expected float drift based on starting location. The blue paddle indicates the starting location and the green paths shows the drift.

the float at each surfacing is shown in blue with the arrows indicating the direction. The red points show where, at each step of re-planning, the float was predicted to travel. Since the re-planning was performed two cycles ahead, two surfacing locations are displayed. As shown, the expected control for neither of the floats performed very well.

This poor performance is not surprising as the current velocities in the ROMS model in the area near both floats was not very accurate, as can be seen in Fig. 16 and Fig. 17. Because EM-APEX is designed to get velocity data, it provides a good opportunity to compare collected data to the ROMS model. The plots show the zonal and meridional currents of the interpolated point in the ROMS model at each location where velocity data was collected by the float. In fact, when comparing these values to the values used in the simulated experiment in Table 3, the RMS errors experienced in the deployment are in the range of the worst models. The RMS errors for the meridional current for both floats were similar to the worst model, the RMS error for the zonal current for float 6667 was similar to the second to worst model, and the RMS error for the zonal current of float 6665 was almost twice as much as the worst model. This can be seen by the lines indicating the RMS errors of the currents for the deployed floats compared to the simulated models in Fig. 10, Fig. 11, Fig. 12, and Fig. 13. Starting in this range the benefits of continuous planning started to decrease and even batch planning scores started to approach those of the deep baseline control sequence.

The poor correlation between the ROMS and the float collected velocities is most likely due to the front that was coming in during the deployment that even caused the deployment to be cut short.




Figure 15: The deployed path and predicted path at each re-planned step in the deployed path for floats 6665 and 6667.




Figure 17: Zonal and Meridional currents found along the path for float 6667 in the ROMS model and actually experienced in the deployment.

Table 4 summarizes the RMS errors and correlation coefficients from the data presented in Fig. 16 and Fig. 17.

<b>6665</b>	<b>RMS Error</b>	<b>Corr Coef</b>	<b>Min Vel (m/s)</b>	<b>Max Vel (m/s)</b>
Zonal	0.24	-0.28	-0.46	0.38
Meridional	0.13	0.42	-0.39	0.39

<b>6667</b>	<b>RMS Error</b>	<b>Corr Coef</b>	<b>Min Vel (m/s)</b>	<b>Max Vel (m/s)</b>
Zonal	0.12	0.33	-0.23	0.53
Meridional	0.13	0.05	-0.42	0.29

Table 4: Root mean square error and correlation coefficient for the deployed velocities collected by the EM-APEX floats compared to the ROMS velocities. The minimum and maximum velocities collected by the floats is also shown.




Figure 16: Zonal and Meridional currents found along the path for float 6665 in the ROMS model and actually experienced in the deployment.

The results from the April 2015 deployment reinforce the thesis of this paper. In cases where the current model provides significant information, the model information can be used improve to float control. In cases where the model provides no or bad information, performance will be poor and even in some cases worse than open loop algorithms such as constant profiling.

## Related Work

Path planning for underwater vehicles has been widely studied; however, most of this work focuses on marine vehicles with greater control such as Sea gliders and Autonomous Underwater Vehicles (AUVs). A notable exception is (Dahl et al. 2011) which examined the problem of optimizing coverage across the oceans for a large number of floats, but only

considered a constant depth and a greedy algorithm. Much more research has been done on glider planning, where there is some control for choosing a direction of motion, but it is less than the current velocity. (Thompson et al. 2010) also uses the ROMS model, but calculates reachability envelopes using wavefront propagation for glider path planning. The work in (Eriksen et al. 2001) describes Seaglider, a glider that is manually controlled from the shore, and is sometimes controlled to maintain position. No ocean model similar to ROMS was used. (Alvarez, Garau, and Caiti 2007) also does not use an ocean model, but instead uses synthetic data with general algorithms to control a set of floats and gliders. Like the work in this paper, (Rao and Williams 2009) uses an A\* graph search algorithm; however, that work assumes that currents change slowly with time and compute the path across many nodes in a single time step, whereas we have many time steps within a single cell. Instead of trying to remain near a specific location, (Pereira et al. 2013) focuses on gliders that are attempting to avoid surfacing in dangerous areas, such as shipping lanes. (Grasso et al. 2010) focuses on the prediction of the glider location, analyzes the accuracy of the predictive model, and uses a physics based control model. Using an asset with more control, (Cashmore et al. 2014) explores the problem of autonomously maneuvering near a site for inspection using an AUV with probabilistic modeling for uncertainty. Autonomous marine vehicles have even been proposed to explore Titan, a moon in the Saturnian system (Pedersen et al. 2015; ESA/NASA 2009; Stofan et al. 2009). (Leonard et al. 2010) present a controls-based methods to guide a set of gliders along coordinated paths. Their approach does not use a model of currents, but does adaptively guides assets back on to paths and desired spacing them along racetracks if the vehicles are perturbed by currents. Therefore this approach would counter-act currents but does not use a projective forward model as our approach does.

Continuous planning has become more prevalent in recent years and the evolution of this planning technique, with respect to multiple assets, is clearly described in (Durfee et al. 1999). (Myers 1999) describes a Continuous Planning and Execution Framework (CPEF), which integrates planning and execution through plan generation, monitoring, execution, and repair. Using an iterative repair process, as well as user interaction, CPEF is able to plan in unpredictable and dynamic environments, which is shown through tests in a simulation of an air-campaign for dominance. (Chien et al. 2000) presents Continuous Activity Scheduling Planning Execution and Replanning (CASPER), which also uses iterative repair as part of continuous planning, specifically for autonomous spacecraft control. (Branch et al. 2016) uses continuous planning to control AUVs, Seagliders, and Wave Gliders, also using different fidelity ROMS models, to follow short distance patterns.

## Future Work

There are many different extensions that are possible from this work. A different objective function that guides the float along a line or a moving point could be developed. Allowing the float to move to any depth at any time, instead of requir-

ing full profiles to the same depth could be tested. Deployment and retrieval cost could be included in the objective. Different assets could be included, such as a glider or AUV, to give more flexibility in the control. Multiple assets of different types could be controlled simultaneously with different objectives and goal locations. New methods for understanding the information gain in the different models could be created as well as using the information about the predictive accuracy in the model to change how the planning is performed. For example, if it is known that the model performs poorly after a particular time, that information could be used to adjust the planning algorithm.

## Conclusions

When performing predictive path planning for underwater vehicles, the model used to represent the ocean always has some limitations in terms of predictive accuracy. This experiment has shown that the amount of knowledge in the planning model used to generate the control sequence for a vertical profiling float attempting to perform virtual mooring affects the benefits of performing planning over simply evenly spacing the profiles of the float. Specifically, as the predictive accuracy of the planning model decreases, represented by the RMS error of the planning model to the nature model, the benefits of planning also decrease.

One method to counteract the poor information in the planning model is to perform continuous planning. We have shown that continuous planning is beneficial when the planning model does not match the nature model, but there is still some valid information in the model used for planning. The RMS error of the model can be used to determine if continuous planning is worthwhile.

## Acknowledgments

Portions of this work were performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

## References

- Alvarez, A.; Garau, B.; and Caiti, A. 2007. Combining networks of drifting profiling floats and gliders for adaptive sampling of the ocean. In *IEEE Intl Conf Robotics and Automation, 2007*, 157–162. IEEE.
- Bluefin Robotics Corporation. “Vehicles, Batteries & Services”. <http://www.bluefinrobotics.com/vehicles-batteries-and-services/>. Accessed Feb, 2016.
- Branch, A.; Troesch, M.; Chu, S.; Chien, S.; Chao, Y.; Farara, J.; and Thompson, A. 2016. Evaluating scientific coverage strategies for a heterogeneous fleet of marine assets using a predictive model of ocean currents. In *Scheduling and Planning Applications Workshop, International Conference on Automated Planning and Scheduling 2016*. ICAPS.
- Cashmore, M.; Fox, M.; Larkworthy, T.; Long, D.; and Magazzeni, D. 2014. Auv mission control via temporal planning. In *IEEE International Conference on Robotics and Automation (ICRA), 2014*, 6535–6541. IEEE.

- Chao, Y.; Li, Z.; Farrara, J.; McWilliams, J. C.; Bellingham, J.; Capet, X.; Chavez, F.; Choi, J.-K.; Davis, R.; Doyle, J.; et al. 2009. Development, implementation and evaluation of a data-assimilative ocean forecasting system off the central california coast. *Deep Sea Research Part II: Topical Studies in Oceanography* 56(3):100–126.
- Chassignet, E. P.; Hurlburt, H. E.; Smedstad, O. M.; Halliwell, G. R.; Hogan, P. J.; Wallcraft, A. J.; Baraille, R.; and Bleck, R. 2007. The hycom (hybrid coordinate ocean model) data assimilative system. *J Marine Systems* 65(1):60–83.
- Chien, S. A.; Knight, R.; Stechert, A.; Sherwood, R.; and Rabideau, G. 2000. Using iterative repair to improve the responsiveness of planning and scheduling. In *Proc Artificial Intelligence Planning Scheduling (AIPS)*, 300–307. AAAI.
- Dahl, K. P.; Thompson, D. R.; McLaren, D.; Chao, Y.; and Chien, S. 2011. Current-sensitive path planning for an underactuated free-floating ocean sensorweb. In *IEEE/RSJ Intl Conf on Intelligent Robots and Systems (IROS)*, 2011, 3140–3146. IEEE.
- Durfee, E. H.; Ortiz Jr, C. L.; Wolverson, M. J.; et al. 1999. A survey of research in distributed, continual planning. *Ai magazine* 20(4):13–22.
- Eriksen, C. C.; Osse, T. J.; Light, R. D.; Wen, T.; Lehman, T. W.; Sabin, P. L.; Ballard, J. W.; and Chiodi, A. M. 2001. Seaglider: A long-range autonomous underwater vehicle for oceanographic research. *Oceanic Engineering, IEEE Journal of* 26(4):424–436.
- ESA/NASA. 2009.
- Farrara, J. D.; Chao, Y.; Zhang, H.; Seegers, B. N.; Teel, E. N.; Caron, D. A.; Howard, M.; Jones, B. H.; Robertson, G.; Rogowski, P.; and Terrill, E. 2015. Oceanographic conditions during the orange county sanitation district diversion experiment as revealed by observations and model simulations. *Estuarine, Coastal and Shelf Science*.
- Grasso, R.; Cecchi, D.; Cococcioni, M.; Trees, C.; Rixen, M.; Alvarez, A.; and Strode, C. 2010. Model based decision support for underwater glider operation monitoring. In *OCEANS 2010*, 1–8. IEEE.
- Jet Propulsion Laboratory. “AirSWOT”. <https://swot.jpl.nasa.gov/airswot/>. Accessed Feb, 2016.
- Jet Propulsion Laboratory. “Earth Science Airborne Program”. <http://airbornescience.jpl.nasa.gov/instruments/airswot>. Accessed Feb, 2016.
- Kongsberg Mairtime AS. “Autonomous Underwater Vehicles - AUV”. [http://www.km.kongsberg.com/ks/web/nokbg0240.nsf/AllWeb/D5682F98CBFBC05AC1257497002976E4?](http://www.km.kongsberg.com/ks/web/nokbg0240.nsf/AllWeb/D5682F98CBFBC05AC1257497002976E4?OpenDocument) OpenDocument. Accessed Feb, 2016.
- Leonard, N. E.; Paley, D. A.; Davis, R. E.; Fratantoni, D. M.; Lekien, F.; and Zhang, F. 2010. Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in monterey bay. *J Field Robotics* 27(6):718–740.
- Li, P.; Chao, Y.; Vu, Q.; Li, Z.; Farrara, J.; Zhang, H.; and Wang, X. 2006. Ourocean-an integrated solution to ocean monitoring and forecasting. In *OCEANS 2006*, 1–6. IEEE.
- Mellor, G. L. 2004. *Users guide for a three dimensional, primitive equation, numerical ocean model*. Princeton, NJ: Princeton University.
- Myers, K. L. 1999. Cpef: A continuous planning and execution framework. *AI Magazine* 20(4):63–69.
- OceanServer Technology, Inc. “Ecomapper AUV”. <http://www.ysisystems.com/productsdetail.php?EcoMapper-Autonomous-Underwater-Vehicle-9>. Accessed Feb, 2016.
- Pedersen, L.; Smith, T.; Lee, S. Y.; and Cabrol, N. 2015. Planetary lakelander - a robotic sentinel to monitor remote lakes. *J Field Robotics* 32(6):860–879.
- Pereira, A. A.; Binney, J.; Hollinger, G. A.; and Sukhatme, G. S. 2013. Risk-aware path planning for autonomous underwater vehicles using predictive ocean models. *J Field Robotics* 30(5):741–762.
- Rao, D., and Williams, S. B. 2009. Large-scale path planning for underwater gliders in ocean currents. In *Australasian Conf Robotics and Automation (ACRA)*, Sydney.
- Robinson, A. R. 1999. Forecasting and simulating coastal ocean processes and variabilities with the Harvard Ocean Prediction System. In Mooers, C. N. K., ed., *Coastal Ocean Prediction*, AGU Coastal and Estuarine Studies Series. Wash., DC: American Geophysical Union. 77–100.
- Sanford, T. B.; Dunlap, J. H.; Carlson, J.; Webb, D. C.; Girtton, J. B.; et al. 2005. Autonomous velocity and density profiler: Em-apex. In *Proceedings of the IEEE/OES Eighth Working Conference on Current Measurement Technology*, 2005, 152–156. IEEE.
- Stofan, E.; Lorenz, R.; Lunine, J.; Aharonson, O.; Bierhaus, E.; Clark, B.; Griffith, C.; Harri, A.-M.; Karkoschka, E.; Kirk, R.; Kantsiper, B.; Mahaffy, P.; Newman, C.; Ravine, M.; Trainer, M.; Waite, H.; and Zarnecki, J. 2009. Titan mare explorer (TiME): first in situ exploration of an extraterrestrial sea.
- Thompson, D. R.; Chien, S.; Chao, Y.; Li, P.; Cahill, B.; Levin, J.; Schofield, O.; Balasuriya, A.; Petillo, S.; Arrott, M.; et al. 2010. Spatiotemporal path planning in strong, dynamic, uncertain currents. In *IEEE Intl Conf on Robotics and Automation (ICRA)*, 2010, 4778–4783. IEEE.
- Troesch, M.; Chien, S.; Chao, Y.; and Farrara, J. 2016. Planning and control of marine floats in the presence of dynamic, uncertain currents. In *Proceedings of ICAPS 2016*. AAAI.
- Wang, X.; Chao, Y.; Thompson, D. R.; Chien, S. A.; Farrara, J.; Li, P.; Vu, Q.; Zhang, H.; Levin, J. C.; and Gangopadhyay, A. 2013. Multi-model ensemble forecasting and glider path planning in the mid-atlantic bight. *Continental Shelf Research* 63:S223–S234.
- Woods Hole Oceanographic Institution. “Floats & Drifters”. <https://www.whoi.edu/main/instruments/floats-drifters>. Accessed Feb, 2016.
- YSI Systems. “IVER Autonomous Underwater Vehicle”. <http://www.iver-auv.com>. Accessed Feb, 2016.